

Better algorithms for satisfiability problems for formulas of bounded rank-width

Robert Ganian, Petr Hliněný, and Jan Obdržálek *

Faculty of Informatics, Masaryk University
Botanická 68a, Brno, Czech Republic

{xganian1,hlineny,obdrzalek}@fi.muni.cz

Abstract. We provide a parameterized polynomial algorithm for the propositional model counting problem #SAT, the runtime of which is single-exponential in the rank-width of a formula. Previously, analogous algorithms have been known – e.g. [Fischer, Makowsky, and Ravve] – with a single-exponential dependency on the clique-width of a formula. Our algorithm thus presents an exponential runtime improvement (since clique-width reaches up to exponentially higher values than rank-width), and can be of practical interest for small values of rank-width. We also provide an algorithm for the MAX-SAT problem along the same lines.

Key words: propositional model counting; satisfiability; rank-width; clique-width; parameterized complexity.

1 Introduction

The satisfiability problem for Boolean formulas in conjunctive normal form (known as SAT) has been of great practical and theoretical interest for decades. It is known to be NP-complete, even though many instances are practically solvable using the various SAT-solvers. We focus on two well-known generalizations of this problem, namely #SAT and MAX-SAT. In #SAT – otherwise known as the *propositional model counting* problem, the goal is to compute the number of satisfying truth assignments for an input formula ϕ , whereas in MAX-SAT we ask for the maximum number of simultaneously satisfiable clauses of ϕ . It is known that computing #SAT is #P-hard [19] and that MAX-SAT is already NP-hard to approximate within some constant [1].

In light of these hardness results, we may ask what happens if we restrict ourselves to some subclass of inputs. The parameterized algorithmics approach is suitable in such a case. Let k be some parameter associated with the input instance. Such a decision problem is said to be *fixed-parameter tractable (FPT)* if it is solvable in time $\mathcal{O}(n^p \cdot f(k))$ for some constant p and a computable function f . So the running time is polynomial in the size n of the input, but can be e.g. exponential in the parameter k . Obviously the specific form of f

* This research has been supported by the Czech research grant GA 201/09/J021 and by the Institute for Theoretical Computer Science ITI, project 1M0545.

plays an important role in practical applicability of any such algorithm – while FPT algorithms with single-exponential f can be feasible for non-trivial values of the parameter, a double-exponential f would make the algorithm impractical for almost all values of k .

But what are suitable parameters for satisfiability problems? In the particular case of MAX-SAT, one can consider the desired number of satisfied or unsatisfied clauses as a parameter of the input, such as in [4, 17], respectively. Although, such approach is not at all suitable for #SAT which is our prime interest in this paper.

Another approach used for instance by Fischer, Makowsky and Ravve [7] represents the formula ϕ as a formula graph F_ϕ (nodes of which are the clauses and variables of ϕ , see Definition 2.7), and exploits the fact that for graphs there are many known (and intensively studied) so called *width parameters*. In [7] the authors presented FPT algorithms for the #SAT problem in the case of two well known width parameters – *tree-width* and *clique-width*. A similar idea was used by Georgiou and Papakonstantinou [9] also for the MAX-SAT problem and by Samer and Szeider [18] for #SAT.

The latter algorithms work by dynamic programming on tree-like decompositions related to the width parameters (tree-decompositions and clique-decompositions – often called k -expressions – in the cases above). However, there is the separate issue of the complexity of computing the width of the formula graph and its decomposition. In the case of tree-width this can be done in FPT [2]. For the much more general clique-width (every graph of bounded tree-width also has bounded clique-width, while the converse does not hold) there exist no such algorithms and we rely on approximations or an oracle. In [18] the authors made the following statement on this issue:

A single-exponential algorithm (for #SAT) is due to Fisher, Makowsky, and Ravve [7]. However, both algorithms rely on clique-width approximation algorithms. The known polynomial-time algorithms for that purpose admit an exponential approximation error [12] and are of limited practical value.

The exponential approximation error mentioned in this statement results by bounding the clique-width by a another, fairly new, width parameter called *rank-width* (Definition 2.1). Rank-width is bounded if and only if clique-width is bounded, but its value can be exponentially lower than that of clique-width (Theorem 2.2 a,b). And since clique-width generalizes tree-width, so does rank-width (Theorem 2.2 c). Moreover, for rank-width we can efficiently compute the related decomposition (Theorem 2.3), which is in stark contrast to the case for clique-width. Therefore an algorithm which is linear in the formula size and *single-exponential in its rank-width* challenges the claim quoted above, and can be of real practical value. In this paper we present such algorithms for the problems #SAT and MAX-SAT. More precisely we prove the following two results:

Theorem 1.1. *Both the #SAT and MAX-SAT problems have FPT algorithms running in time*

$$\mathcal{O}(t^3 \cdot 2^{3t(t+1)/2} \cdot |\phi|)$$

where t is the rank-width of the input instance (CNF formula) ϕ .

We refer to further Theorems 2.11, 3.1 and 4.1 for details.

Note that our results present an *actual exponential runtime improvement* in the parameter over any algorithm utilizing the clique-width measure, including aforementioned [7]. This is since any parameterized algorithm \mathcal{A} for a SAT problem has to depend at least exponentially on the clique-width of a formula (unless the exponential time hypothesis fails), and considering typical instances ϕ as from Proposition 2.10 b, such an algorithm \mathcal{A} then runs in time doubly-exponential in the rank-width of ϕ .

As for potential practical usefulness of Theorem 1.1, note that there are no “large constants” hidden in the \mathcal{O} -notation. One may also ask whether there are any interesting classes of graphs of low rank-width. The answer is a resounding YES, since already for $t = 1$ we obtain the very rich class of distance-hereditary graphs. Rank-width indeed is a very general graph width measure.

The approach we use to prove both parts of Theorem 1.1 quite naturally extends the clever and skilled new algebraic methods of designing parameterized algorithms for graphs of bounded rank-width, e.g. [6, 3, 8], to the area of SAT problems. Yet, this is not a trivial extension—we remark that a straightforward translation of the algorithm of [7] from clique-width expressions to rank-decompositions (which is easily possible) would result just in a double-exponential runtime dependency on the rank-width.

The rest of the paper is organized as follows: In Section 2 we present the rank-width measure and some related technical considerations. This is applied to signed graphs of SAT formulas. Section 3 then presents our FPT algorithm for the #SAT problem (Theorem 3.1 and Algorithm 3.6), and Section 4 the similar algorithm for MAX-SAT (Theorem 4.1). We conclude with some related observations.

2 Overview of the rank-width measure

Graph rank-width [16], the core concept of our paper, is not so well known, and hence we give a detailed technical introduction to this concept and its application to CNF formulas in this section. Readers familiar with the concept of rank-width (and parse trees for rank-width) may proceed directly to Section 2.3.

2.1 Branch-width and rank-width

The usual way of defining rank-width is via the branch-width of the cut-rank function (Definition 2.1). A set function $f : 2^M \rightarrow \mathbb{Z}$ is *symmetric* if $f(X) = f(M \setminus X)$ for all $X \subseteq M$. A tree is *subcubic* if all its nodes have degree at

most 3. For a symmetric function $f : 2^M \rightarrow \mathbb{Z}$ on a finite ground set M , the branch-width of f is defined as follows:

A *branch-decomposition* of f is a pair (T, μ) of a subcubic tree T and a bijective function $\mu : M \rightarrow \{t : t \text{ is a leaf of } T\}$. For an edge e of T , the connected components of $T \setminus e$ induce a bipartition (X, Y) of the set of leaves of T . The *width* of an edge e of a branch-decomposition (T, μ) is $f(\mu^{-1}(X))$. The *width* of (T, μ) is the maximum width over all edges of T . The *branch-width* of f is the minimum of the width of all branch-decompositions of f .

Definition 2.1 (Rank-width [16]). For a simple graph G and $U, W \subseteq V(G)$, let $\mathbf{A}_G[U, W]$ be the matrix defined over the two-element field $\text{GF}(2)$ as follows: the entry $a_{u,w}$, $u \in U$ and $w \in W$, of $\mathbf{A}_G[U, W]$ is 1 if and only if uw is an edge of G . The *cut-rank* function $\rho_G(U) = \rho_G(W)$ then equals the rank of $\mathbf{A}_G[U, W]$ over $\text{GF}(2)$ where $W = V(G) \setminus U$. A *rank-decomposition* (see Figure 1) and *rank-width* of a graph G is the branch-decomposition and branch-width of the cut-rank function ρ_G of G on $M = V(G)$, respectively.

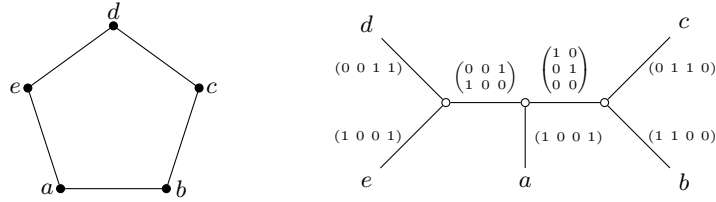


Fig. 1. A rank-decomposition of the graph cycle C_5 , showing the matrices involved in evaluation of its cut-rank function on the edges of the decomposition.

As already mentioned in the introduction, rank-width is closely related to clique-width and more general than better known tree-width. Indeed:

Theorem 2.2. Let G be a simple graph, and $\text{tw}(G)$, $\text{bw}(G)$, $\text{cwd}(G)$, $\text{rwd}(G)$ denote in this order the tree-width, branch-width, clique-width, and rank-width of G . Then the following hold

- a) [16] $\text{rwd}(G) \leq \text{cwd}(G) \leq 2^{\text{rwd}(G)+1} - 1$,
- b) [5] the clique-width $\text{cwd}(G)$ can reach up to $2^{\text{rwd}(G)/2-1}$,
- c) [15] $\text{rwd}(G) \leq \text{bw}(G) \leq \text{tw}(G) + 1$,
- d) [folklore] $\text{tw}(G)$ cannot be bounded from above by $\text{rwd}(G)$, e.g. the complete graphs have rank-width 1 while their tree-width is unbounded,
- e) [14] $\text{rwd}(G) = 1$ if and only if G is a distance-hereditary graph.

Although rank-width and clique-width are “tied together” (a), one of the crucial advantages of rank-width is its parameterized tractability (on the other hand, it is not known how to efficiently test $\text{cwd}(G) \leq k$ for $k > 3$):

Theorem 2.3 ([12]). *There is an FPT algorithm that, for a fixed parameter t and a given graph G , either finds a rank-decomposition of G of width at most t or confirms that the rank-width of G is more than t .*

2.2 Labeling parse trees for rank-width

Unlike for tree-width and clique-width, the standard definition of rank-decompositions is not suitable for the immediate design of efficient algorithms. To this end, closely following Courcelle and Kanté [6], we have introduced so-called *labeling parse trees* [8] (Definition 2.5 and Figure 2) – a powerful formalism for dynamic programming design on graphs of bounded rank-width. The basic idea is to transform rank-decompositions into suitable parse trees and have algorithms use them instead of the decomposition.

A t -labeling of a graph is a mapping $lab : V(G) \rightarrow 2^{[t]}$ where $[t] = \{1, 2, \dots, t\}$ is the set of *labels*. Having a graph G with an associated t -labeling lab , we refer to the pair (G, lab) as to a t -labeled graph and use notation \bar{G} . We will often view a t -labeling of G equivalently as a mapping $V(G) \rightarrow \text{GF}(2)^t$ to the *binary vector space* of dimension t , where $\text{GF}(2)$ is the two-element finite field.

Definition 2.4. Considering t -labeled graphs $\bar{G}_1 = (G_1, lab^1)$ and $\bar{G}_2 = (G_2, lab^2)$, a t -labeling join $\bar{G}_1 \otimes \bar{G}_2$ is defined on the disjoint union of G_1 and G_2 by adding all edges (u, v) such that $|lab^1(u) \cap lab^2(v)|$ is odd, where $u \in V(G_1), v \in V(G_2)$. (Alternatively, $\{u, v\}$ is an edge of $\bar{G}_1 \otimes \bar{G}_2$ if and only if $lab^1(u) \cdot lab^2(v) = 1$ over $\text{GF}(2)$.) The resulting graph is unlabeled.

A t -relabeling is a mapping $f : [t] \rightarrow 2^{[t]}$. In linear algebra terms, a t -relabeling f is in a natural one-to-one correspondence with a *linear transformation* $f : \text{GF}(2)^t \rightarrow \text{GF}(2)^t$, i.e. a $t \times t$ binary matrix \mathbf{R}_f . For a t -labeled graph $\bar{G} = (G, lab)$ we define $f(\bar{G})$ as the same graph with a vertex t -labeling $lab' = f \circ lab$. Here $f \circ lab$ stands for the linear transformation f applied to the labeling lab , or equivalently $lab' = lab \times \mathbf{R}_f$ as matrix multiplication over $\text{GF}(2)^t$.

Definition 2.5 (Labeling parse tree [6, 8]). Let \odot be a nullary operator creating a single new graph vertex of label $\{1\}$. For t -relabelings $f_1, f_2, g : [t] \rightarrow 2^{[t]}$, let $\otimes[g | f_1, f_2]$ be a binary operator—called *t -labeling composition*—over pairs of t -labeled graphs $\bar{G}_1 = (G_1, lab^1)$ and $\bar{G}_2 = (G_2, lab^2)$ defined (cf. 2.4)

$$\bar{G}_1 \otimes[g | f_1, f_2] \bar{G}_2 = \bar{H} = (\bar{G}_1 \otimes g(\bar{G}_2), lab)$$

where the new labeling is $lab(v) = f_i \circ lab^i(v)$ for $v \in V(G_i)$, $i = 1, 2$. In other words, $\{u, v\} \in E(H)$ where $u \in V(G_1)$, $v \in V(G_2)$, if and only if $lab^1(u) \times \mathbf{R}_g^T \times lab^2(v)^T = 1$ over $\text{GF}(2)$ (cf. Courcelle and Kanté [6]).

A t -labeling parse tree T is a finite rooted ordered subcubic tree (with the root degree at most 2) such that

- all leaves of T contain the \odot symbol, and
- each internal node of T contains one of the t -labeling composition symbols.

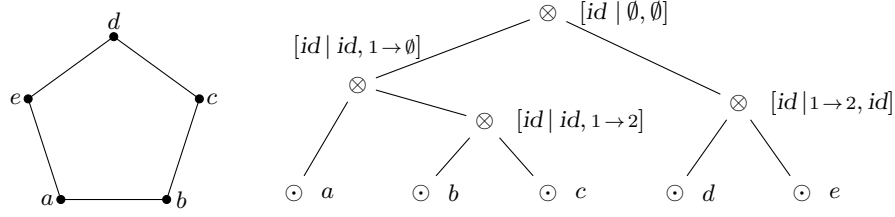


Fig. 2. An example of a 2-labeling parse tree which generates a cycle C_5 , with symbolic relabelings at the nodes (id denotes the relabeling preserving all labels, and \emptyset is the relabeling “forgetting” all labels).

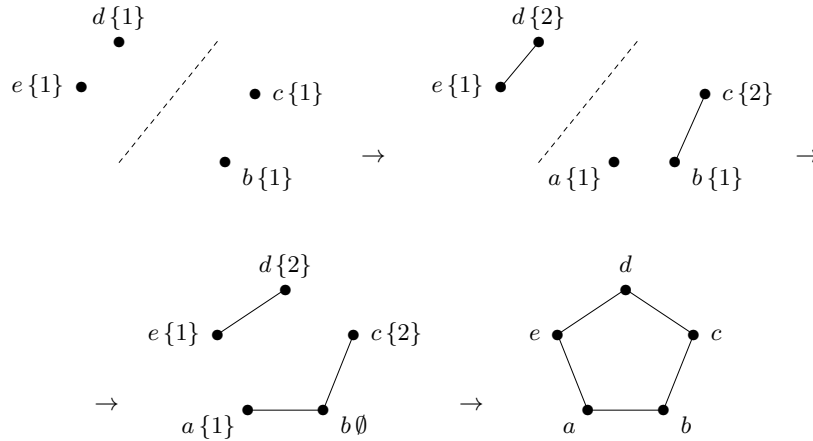


Fig. 3. “Bottom-up” generation of C_5 by the parse tree from Fig. 2.

A parse tree T then *generates* (parses) the graph G which is obtained by successive leaves-to-root applications of the operators in the nodes of T .

See Figures 2 and 3. The crucial statement is that rank-decompositions are exactly equivalent to labeling parse trees:

Theorem 2.6 ([6, 8]). *A graph G has rank-width at most t if and only if (some labeling of) G can be generated by a t -labeling parse tree. Furthermore, a width- t rank-decomposition of G can be transformed into a t -labeling parse tree on $\Theta(|V(G)|)$ nodes in time $\mathcal{O}(t^2 \cdot |V(G)|^2)$.*

2.3 Signed graphs and rank-width of CNF formulas

Although there are several methods for converting formulas to graphs, the most common and perhaps most natural approach uses so-called *signed graphs* (e.g. [7, 18, 11]). A signed graph is a graph G with two edge sets $E^+(G)$ and $E^-(G)$. We refer to its respective positive and negative subgraphs as to G^+ and G^- . Notice that G^+ and G^- are edge-disjoint and $G = G^+ \cup G^-$.

Definition 2.7. The signed graph F_ϕ of a CNF formula ϕ is defined as follows:

- $V(F_\phi) = W \cup C$ where W is the set of variables occurring in ϕ and C is the set of clauses of ϕ .
- For $w \in W$ and $c \in C$, it is $wc \in E^+(F_\phi)$ iff the literal ‘ w ’ occurs in c .
- For $w \in W$ and $c \in C$, it is $wc \in E^-(F_\phi)$ iff the literal ‘ $\neg w$ ’ occurs in c .

Since signed graphs have two distinct edge sets, the definition of rank-width needs to be modified to reflect this. It should be noted that simply using two separate, independent decompositions would not work – the bottom-up dynamic programming algorithm we are going to use will need information from both edge sets at every node to work properly. Instead, one may define, analogically to Definition 2.1, the *signed rank-width* of a signed graph G as the branch-width of the signed cut-rank function $\rho_G^\pm(U) = \rho_{G^+}(U) + \rho_{G^-}(U)$.

Definition 2.8 (Rank-width of formulas). The *(signed) rank-width* $\text{rwd}(\phi)$ of a CNF formula ϕ is the signed rank-width of the signed formula graph F_ϕ .

Although our signed rank-width is essentially equivalent to an existing concept of bi-rank-width of directed graphs as introduced by Kanté [13] (in the bipartite case, at least), the latter concept is not widely known and its introduction in the context of CNF formulas would bring only additional technical complications. In the SAT context, it is more natural and easier to deal with undirected signed graphs. Hence we introduce, following previous Definition 2.5, (t^+, t^-) -labeling parse trees which will be equivalent to signed rank-width (up to a factor of 2, see Theorem 2.11) in a way analogical to Theorem 2.6.

Definition 2.9. A (t^+, t^-) -labeling parse tree $T = (T^+, T^-)$ of a signed graph G is a pair (T^+, T^-) of two labeling parse trees T^+ and T^- such that:

- I. T^+ (T^-) is a t^+ -labeling (t^- -labeling) parse tree generating G^+ (G^-), and
- II. The underlying rooted ordered trees of T^+ and T^- are identical.

With a slight abuse of terminology, we will refer to the pair of subtrees of T^+ and T^- rooted at a common node s as to a subtree of T rooted at s .

Analogically to labeled graphs of Section 2.2, we call a signed graph G with associated pair of labelings $\text{lab}^+ : V(G) \rightarrow 2^{[t^+]}$ and $\text{lab}^- : V(G) \rightarrow 2^{[t^-]}$ a (t^+, t^-) -labeled graph $\tilde{G} = (G, \text{lab}^+, \text{lab}^-)$. We shortly refer to the t^+ -labeled graph (G^+, lab^+) as to \tilde{G}^+ , and analogically to \tilde{G}^- . The scope of the join operation \otimes (Definition 2.4) can then be extended in a natural way as

$$\tilde{G}_1 \otimes \tilde{G}_2 = (\tilde{G}_1^+ \otimes \tilde{G}_2^+) \cup (\tilde{G}_1^- \otimes \tilde{G}_2^-).$$

In our paper we propose signed rank-width as a way of measuring complexity of formulas that fares significantly better than previously considered signed clique-width of F_ϕ (e.g. [7]). *Signed clique-width* is the natural extension of clique-width having two separate operators for creating the ‘plus’ and the ‘minus’ edges. The advantage of our approach is witnessed by the following two claims.

Proposition 2.10. *Let ϕ be an arbitrary CNF formula and F_ϕ its signed graph of signed clique-width $\text{cwd}(\phi)$. Then the following are true*

- a) $\text{rwd}(\phi) \leq 2\text{cwd}(\phi)$,
- b) *there exist instances ϕ such that $\text{cwd}(\phi) \geq 2^{\text{rwd}(\phi)/4-1}$.*

Proof. a) Assume a signed k -expression tree S for ϕ where $k = \text{cwd}(\phi)$. Clearly, S gives ordinary k -expression trees for each of F_ϕ^+ , F_ϕ^- . Now, analogically to Theorem 2.2(a) [16], the rank-decompositions of F_ϕ^+ and F_ϕ^- with the same underlying tree as S have widths $\leq k$ each, and hence $\text{rwd}(\phi) \leq k+k = 2\text{cwd}(\phi)$.

b) We define ϕ such that $F_\phi^+ = G$ where, cf. Theorem 2.2(b), $\text{cwd}(G) \geq 2^{\text{rwd}(G)/2-1}$, and F_ϕ^- is arbitrary such that its rank-decomposition inherited from that of G has width $\leq \text{rwd}(G)$. Then $\text{rwd}(\phi) \leq 2\text{rwd}(G)$ and the claim follows since $\text{cwd}(\phi) \geq \text{cwd}(G)$. \square

Theorem 2.11 ([12]). *There is an FPT algorithm that, for a fixed parameter t and a given CNF formula ϕ , either finds a (t^+, t^-) -labeling parse tree for the formula graph F_ϕ where $t^+ \leq t$ and $t^- \leq t$, or confirms that the signed rank-width of ϕ is more than t .*

Proof. We show that the algorithm of [12] (Theorem 2.3) can be used to compute also the signed rank-width $\leq t$ of a signed graph $G = F_\phi$. Indeed, we define a new graph G' as the union of G^+ and a vertex-disjoint copy G_0^- of G^- , and a partition \mathcal{P} of $V(G')$ as the collection of the corresponding vertex pairs from $V(G^+) \times V(G_0^-)$. Then we call the algorithm of [12] to either compute a \mathcal{P} -partitioned rank-decomposition of G' of width $\leq t$, or confirm that the \mathcal{P} -partitioned width is $> t$. This width is exactly our signed rank-width of G since, for any bipartition (U, W) of $V(G')$ not crossing \mathcal{P} , the cut-rank $\rho_{G'}(U)$ of U in G' trivially equals our signed cut-rank $\rho_G^\pm(X) = \rho_{G^+}(X) + \rho_{G^-}(X)$ of $X = U \cap V(G)$. Lastly, applying Theorem 2.6, we separately transform the two “inherited” rank-decompositions of G^+ and G_0^- into parse trees T^+ and T^- , and output together the (t^+, t^-) -labeling parse tree $T = (T^+, T^-)$ of F_ϕ . \square

3 Algorithm for propositional model counting #SAT

This section proves our most important result – the #SAT part of Theorem 1.1.

We remind the readers that the previous best algorithm [7] for #SAT on graphs of bounded clique-width has had a single-exponential runtime dependency on the signed clique-width of a formula (and this dependency cannot be further improved unless the so called exponential time hypothesis fails). Hence by Proposition 2.10.b the worst case scenario for the algorithm of [7] would lead to a double-exponential runtime dependency on the signed rank-width of the formula. On the other hand:

Theorem 3.1. *Given a CNF formula ϕ and a (t^+, t^-) -labeling parse tree (Theorem 2.11) of the formula graph F_ϕ (Definition 2.9), there is an algorithm that counts the number of satisfying assignments of ϕ in time*

$$\mathcal{O}(t^3 \cdot 2^{3t(t+1)/2} \cdot |\phi|) \quad \text{where } t = \max(t^+, t^-).$$

3.1 Informal notes

Our algorithm (see Algorithm 3.6) proving Theorem 3.1 applies the dynamic programming paradigm on the parse trees of the formula graph F_ϕ (constructed by Theorem 2.11). This is, on one hand, a standard approach utilized also by Fischer, Makowsky and Ravve [7]. On the other hand, however, comparing to [7] we achieve an exponential runtime speedup in terms of rank-width. This significant improvement has two main sources (necessary on both sides):

- We heavily apply the basic calculus and tools of linear algebra in the algorithm (which is indeed natural in view of the algebraic definition of rank-width). See the details in Subsection 3.2.
- Our dynamic programming algorithm is built upon the idea of an “expectation” (when processing the parse tree of the input) – in addition to the information recorded about a partial solution processed so far, we also record what is expected from a complementary partial solution coming from the unprocessed part of the input.

Especially the second point deserves an informal explanation before giving a formal description in Definition 3.4. II. The background idea is that the amount of information one has to remember about a partial solution shrinks a lot if one “knows” what the complete solution will look like. Such saving sometimes largely exceeds the cost of keeping an exhaustive list of all possible future “shapes” of complete solutions. This is also our case where the application is quite natural – we may exhaustively preprocess the values of some variables in advance.

The idea of using an “expectation” to speed up a dynamic programming algorithm on a rank-decomposition has first appeared in Bui-Xuan, Telle and Vatshelle [3] in relation to solving the dominating set on graphs of bounded rank-width. This concept has been subsequently formalized and generalized by the authors in [8] (in the so called PCE scheme formalism). Furthermore, it has also been shown [8, Proposition 5.1] that use of the “expectation” concept is unavoidable to achieve speed up for the dominating set problem.

Unfortunately, we cannot simply refer the formalism of [8] here since it was designed for optimization, and not enumeration, problems. We thus have to describe it again from scratch in our Algorithm 3.6.

3.2 Supplementary technical concepts

This part describes several technical concepts needed to formulate all details of coming Algorithm 3.6. It may be skipped during the first reading.

A useful algebraic concept is that of orthogonality. We say that labeling ℓ is *orthogonal* to a set of labelings X if ℓ has even intersection with every element of X (i.e. the scalar product of the labeling vectors is 0 over $\text{GF}(2)$). Remember that for t -labeling parse trees, in order for two vertices become adjacent by the join operation \otimes , their labelings need to have odd intersection, i.e. to be non-orthogonal. The power of orthogonality comes from the following rather trivial claim occurring already in [3, 8]:

Lemma 3.2. Assume t -labeled graphs \bar{G} and \bar{H} , and arbitrary $X \subseteq V(\bar{G})$ and $y \in V(\bar{H})$. In the join graph $\bar{G} \otimes \bar{H}$, the vertex y is adjacent to some vertex in X if and only if the vector subspace spanned by the \bar{G} -labelings of the vertices of X is not orthogonal to the \bar{H} -labeling vector of y in $\text{GF}(2)^t$.

In view of Lemma 3.2, the following result will be useful in deriving the complexity of our algorithm.

Lemma 3.3 ([10], cf. [8, Proposition 6.1]). The number $S(t)$ of subspaces of the binary vector space $\text{GF}(2)^t$ satisfies $S(t) \leq 2^{t(t+1)/4}$ for all $t \geq 12$.

We recall from Definition 2.7 the signed graph F_ϕ of a formula ϕ on a vertex set $V(F_\phi) = W \cup C$ where W is the set of variables and C is the set of clauses of ϕ . An *assignment* is then a mapping $\nu : W \rightarrow \{0, 1\}$. In the course of computation of our algorithm we will need to remember some local information about all satisfying assignments for ϕ . The information to be remembered for each such assignment is formally described by the following definition.

Definition 3.4. Consider an arbitrary (t^+, t^-) -labeling $\tilde{F}_1 = (F_1, \text{lab}^+, \text{lab}^-)$ of a signed subgraph $F_1 \subseteq F_\phi$, and any partial assignment $\nu_1 : V(F_1) \cap W \rightarrow \{0, 1\}$.

We say that ν_1 is an *assignment of shape* $(\Sigma^+, \Sigma^-, \Pi^+, \Pi^-)$ in \tilde{F}_1 if

- I. Σ^+ is the subspace of $\text{GF}(2)^t$ generated by the label vectors $\text{lab}^+(\nu_1^{-1}(1))$ and Σ^- is the subspace of $\text{GF}(2)^t$ generated by $\text{lab}^-(\nu_1^{-1}(0))$, and
- II. Π^+, Π^- are subspaces of $\text{GF}(2)^t$ such that, for every clause $c \in V(F_1) \cap C$, at least one of the following is true
 - c is adjacent to some vertex from $\nu_1^{-1}(1)$ in F_1^+ or to some vertex from $\nu_1^{-1}(0)$ in F_1^- , or
 - the label vector $\text{lab}^+(c)$ is not orthogonal to Π^+ or $\text{lab}^-(c)$ is not orthogonal to Π^- (cf. Lemma 3.2).

Very informally saying, I. states which true literals in F_1 (w.r.t. ν_1) are available to satisfy clauses of F_ϕ , and II. stipulates that every clause in F_1 is satisfied by a true literal in F_1 or is expected to be satisfied by some literal in $F_\phi - V(F_1)$. Note that one partial assignment ν_1 could be of several distinct shapes, which differ in Π^+, Π^- . (This is true even for complete assignments.) Moreover, there is no requirement on Π^+ and Π^- to have an empty intersection with Σ^+, Σ^- and each other. The trivial useful properties of assignment shapes are:

Proposition 3.5. We consider a CNF formula ϕ with the variable set W , and any assignment $\nu : W \rightarrow \{0, 1\}$. Assume \tilde{F}_1, \tilde{F}_2 are (t^+, t^-) -labeled graphs such that $F_\phi = \tilde{F}_1 \otimes \tilde{F}_2$, and let ν_1, ν_2 denote the restrictions of ν to \tilde{F}_1, \tilde{F}_2 .

- a) The assignment ν is satisfying for ϕ if, and only if, there exist subspaces $\Sigma^+, \Sigma^-, \Pi^+, \Pi^-$ of $\text{GF}(2)^t$ such that ν_1 is of shape $(\Sigma^+, \Sigma^-, \Pi^+, \Pi^-)$ in \tilde{F}_1 and ν_2 is of shape $(\Pi^+, \Pi^-, \Sigma^+, \Sigma^-)$ in \tilde{F}_2 .
- b) If, in \tilde{F}_1 , ν_1 is of shape $(\Sigma_0^+, \Sigma_0^-, \Pi_0^+, \Pi_0^-)$ and, at the same time, ν_1 is of shape $(\Sigma_1^+, \Sigma_1^-, \Pi_1^+, \Pi_1^-)$, then $\Sigma_0^+ = \Sigma_1^+$ and $\Sigma_0^- = \Sigma_1^-$.
- c) The assignment ν_1 is satisfying for ϕ_1 – the subformula of ϕ represented by F_1 if, and only if, ν_1 is of shape $(\Sigma^+, \Sigma^-, \emptyset, \emptyset)$ for some subspaces Σ^+, Σ^- .

□

3.3 The dynamic processing algorithm

We now return to our Theorem 3.1, considering a (t^+, t^-) -labeling parse tree T_ϕ of a given formula graph F_ϕ . The core of our bottom-up dynamic processing of T_ϕ is as follows: At every node z such that the subtree of T_ϕ rooted at z parses a (t^+, t^-) -labeled graph \tilde{F}_z , we record an integer-valued array $Table_z$ indexed by all the quadruples of subspaces of $\text{GF}(2)^t$, where $t = \max(t^+, t^-)$. The value of the entry $Table_z[\Sigma^+, \Sigma^-, \Pi^+, \Pi^-]$ is equal to the number of variable assignments in $F_z \subseteq F_\phi$ that are of the shape $(\Sigma^+, \Sigma^-, \Pi^+, \Pi^-)$ in \tilde{F}_z (cf. Definition 3.4).

For a subset $X \subseteq \text{GF}(2)^t$, let $\langle X \rangle$ denote the vector subspace of $\text{GF}(2)^t$ spanned by the points of X . If f is a relabeling, i.e. a linear transformation defined by a binary matrix \mathbf{R}_f , then $f(X)$ denotes the image of X under f , and $f^T(X)$ denotes the image of X under the *transposed* relabeling given by \mathbf{R}_f^T .

Algorithm 3.6 (Theorem 3.1). Given is a CNF formula ϕ and a signed (t^+, t^-) -labeling parse tree T_ϕ of the formula graph F_ϕ .

1. We initialize all entries of $Table_z$ for $z \in V(T_\phi)$ to 0.
2. We process all nodes of T_ϕ in the leaves-to-root order as follows.
 - a) At a clause leaf c of T_ϕ , we set $Table_c[\emptyset, \emptyset, \Pi^+, \Pi^-] \leftarrow 1$ for all subspaces Π^+, Π^- such that at least one of them is not orthogonal to the label vector of $\{1\}$ (and \emptyset stands for the zero subspace).
 - b) At a variable leaf ℓ of T_ϕ , we set $Table_\ell[\langle\{1\}\rangle, \emptyset, \Pi^+, \Pi^-] \leftarrow 1$ and $Table_\ell[\emptyset, \langle\{1\}\rangle, \Pi^+, \Pi^-] \leftarrow 1$ for all pairs Π^+, Π^- .
 - c) Consider an internal node z of T_ϕ , with the left son x and the right son y such that $Table_x$ and $Table_y$ have already been computed.
 - Let the composition operators at z in the labeling parse trees T_ϕ^+, T_ϕ^- be as $\tilde{F}_z^+ = \tilde{F}_x^+ \otimes [g^+ | f_1^+, f_2^+] \tilde{F}_y^+$ and $\tilde{F}_z^- = \tilde{F}_x^- \otimes [g^- | f_1^-, f_2^-] \tilde{F}_y^-$ (cf. Definition 2.9 for T_ϕ).
 - We loop exhaustively over all indices to $Table_x, Table_y, Table_z$, i.e. over all 12-tuples of subspaces $\Sigma_x^+, \Sigma_x^-, \Pi_x^+, \Pi_x^-, \Sigma_y^+, \Sigma_y^-, \Pi_y^+, \Pi_y^-, \Sigma_z^+, \Sigma_z^-, \Pi_z^+, \Pi_z^-$ of $\text{GF}(2)^t$. If all the following are true
$$\begin{aligned} \Sigma_z^+ &= \langle f_1^+(\Sigma_x^+) \cup f_2^+(\Sigma_y^+) \rangle \text{ and } \Sigma_z^- = \langle f_1^-(\Sigma_x^-) \cup f_2^-(\Sigma_y^-) \rangle, \\ \Pi_x^+ &= \langle f_1^{+T}(\Pi_z^+) \cup g^+(\Sigma_y^+) \rangle \text{ and } \Pi_x^- = \langle f_1^{-T}(\Pi_z^-) \cup g^-(\Sigma_y^-) \rangle, \\ \Pi_y^+ &= \langle f_2^{+T}(\Pi_z^+) \cup g^{+T}(\Sigma_x^+) \rangle \text{ and } \Pi_y^- = \langle f_2^{-T}(\Pi_z^-) \cup g^{-T}(\Sigma_x^-) \rangle, \end{aligned}$$
then we add the product $Table_x[\Sigma_x^+, \Sigma_x^-, \Pi_x^+, \Pi_x^-] \cdot Table_y[\Sigma_y^+, \Sigma_y^-, \Pi_y^+, \Pi_y^-]$ to the table entry $Table_z[\Sigma_z^+, \Sigma_z^-, \Pi_z^+, \Pi_z^-]$.
3. We sum up all the entries $Table_r[\Sigma^+, \Sigma^-, \emptyset, \emptyset]$ where r is the root of T_ϕ and Σ^+, Σ^- are arbitrary subspaces of $\text{GF}(2)^t$. This is the resulting number of satisfying assignments of ϕ .

Proof (Algorithm 3.6 / Theorem 3.1). The task is to prove that the computed value $Table_z[\Sigma^+, \Sigma^-, \Pi^+, \Pi^-]$ is indeed equal to the number of assignments in \tilde{F}_z that are of the shape $(\Sigma^+, \Sigma^-, \Pi^+, \Pi^-)$. This is done by structural induction on z ranging from the leaves of T_ϕ to its root. Then, in step 3 of our algorithm, the computed number of satisfying assignments of ϕ is correct by Proposition 3.5 b,c.

- 2a.** If $z = c$ where c is a clause leaf, then F_c defines a formula with one empty (so far false) clause c . There is only one possible assignment in \tilde{F}_c . In order to satisfy c , its labeling $lab^+(c) = lab^-(c) = \{1\}$ should not be orthogonal to expected Π^+ or Π^- (Definition 3.4.II), as done in step 2a.
- 2b.** If $z = \ell$ where ℓ is a variable leaf, then F_ℓ defines a formula with one variable and no clause. There are two assignments of ℓ and no requirement on Π^+, Π^- from Definition 3.4.II. Hence these two assignments contribute 1 each to all the indicated table entries by Definition 3.4.I.
- 2c.** This is the hard core of our proof. By induction both $Table_x, Table_y$ already contain the correct values. Assume we have a partial assignment ν in \tilde{F}_z of shape $(\Sigma_z^+, \Sigma_z^-, \Pi_z^+, \Pi_z^-)$. Then ν defines partial assignments ν_x in \tilde{F}_x and ν_y in \tilde{F}_y which, in turn, uniquely determine the corresponding subspaces Σ_x^+, Σ_x^- and Σ_y^+, Σ_y^- by Proposition 3.5b). It follows from Definition 3.4.II of a shape and from Definition 2.5 of the composition operators $\otimes[g^+ | f_1^+, f_2^+]$ and $\otimes[g^- | f_1^-, f_2^-]$ at z , that ν_x is of shape $(\Sigma_x^+, \Sigma_x^-, \Pi_x^+, \Pi_x^-)$ for some Π_x^+, Π_x^- if
- i. for every clause c not adjacent to $\nu_x^{-1}(1)$ in F_x^+ , nor to $\nu_x^{-1}(0)$ in F_x^- , we have that $f_1^+(lab_x^+(c))$ is not orthogonal to Π_z^+ or $f_1^-(lab_x^-(c))$ is not orthogonal to Π_z^- (informally, c will be satisfied by the expectation at z after relabeling),
 - ii. or $lab_x^+(c)$ is not orthogonal to $g^+(\Sigma_y^+)$ or $lab_x^-(c)$ is not orthogonal to $g^-(\Sigma_y^-)$ (informally, c is satisfied by a true literal coming from \tilde{F}_y in the labeling composition at z).

Note that, e.g., $f_1^+(lab_x^+(c))$ is not orthogonal to $\mathbf{v} \in \Pi_z^+$ iff

$$1 = f_1^+(lab_x^+(c)) \times \mathbf{v}^T = \left(lab_x^+(c) \times \mathbf{R}_{f_1^+} \right) \times \mathbf{v}^T = lab_x^+(c) \times \left(\mathbf{v} \times \mathbf{R}_{f_1^+}^T \right)^T.$$

Hence putting the two disjoint alternatives i,ii for c together, we see that $lab_x^+(c)$ should not be orthogonal to $f_1^{+T}(\Pi_z^+) \cup g^+(\Sigma_y^+)$ or $lab_x^-(c)$ should not be orthogonal to $f_1^{-T}(\Pi_z^-) \cup g^-(\Sigma_y^-)$. This exactly corresponds to the condition on Π_x^+ and Π_x^- in 2c. Analogical fact is true for Π_y^+ and Π_y^- . Therefore, ν_x and ν_y have been accounted for in $Table_x$ and $Table_y$, respectively, and so ν is now counted in $Table_z[\Sigma_z^+, \Sigma_z^-, \Pi_z^+, \Pi_z^-]$.

On the other hand, we have to prove that no assignment is counted more than once in one particular entry $Table_z[\Sigma_z^+, \Sigma_z^-, \Pi_z^+, \Pi_z^-]$. This is not immediate due to a (limited) freedom in a choice of the “expectation” part of shape in the previous arguments. For $Table_x, Table_y$ this claim is true by induction. Any particular partial assignment ν in \tilde{F}_z uniquely determines Σ_z^+, Σ_z^- , and $\Sigma_x^+, \Sigma_x^-, \Sigma_y^+, \Sigma_y^-$ as above. Then the conditions in 2c of the algorithm also uniquely determine Π_x^+, Π_x^- and Π_y^+, Π_y^- (and so their entries in $Table_x, Table_y$). Hence the assignment ν is counted at most once for every particular choice of Π_z^+, Π_z^- , too.

Lastly, we analyze the runtime of our algorithm. Let $S(t)$ be the number of subspaces of $\text{GF}(2)^t$. Every single call to one of the steps 1, 2a, 2b, and 3 of Algorithm 3.6 is proportional to the size of the table which is $\mathcal{O}(S(t)^4)$. One call to 2c in this algorithm actually has to loop over all 6-tuples $\Sigma_x^+, \Sigma_x^-, \Sigma_y^+, \Sigma_y^-, \Pi_z^+, \Pi_z^-$ of subspaces of $\text{GF}(2)^t$, while the remaining 6 subspaces $\Pi_x^+, \Pi_x^-, \Pi_y^+, \Pi_y^-, \Sigma_z^+, \Sigma_z^-$ can be computed in time $\mathcal{O}(t^3)$ each using standard algorithms of linear algebra. Hence this point takes time $\mathcal{O}(t^3 \cdot S(t)^6)$.

For the sake of completeness, we note that there exists [8, Lemma 6.3] an efficient indexing scheme for all the subspaces of $\text{GF}(2)^t$ with query time $\mathcal{O}(t^3)$. Such a scheme can be built in time $\mathcal{O}(2^{3t(t+1)/4} \cdot t^3)$.

Altogether, using Lemma 3.3, our Algorithm 3.6 takes time

$$\mathcal{O}(|V(T_\phi)| \cdot t^3 \cdot S(t)^6) = \mathcal{O}(|V(T_\phi)| \cdot t^3 \cdot 2^{6t(t+1)/4}) = \mathcal{O}(|\phi| \cdot t^3 \cdot 2^{3t(t+1)/2}). \quad \square$$

4 Algorithm for the Max-SAT problem

The same ideas as presented in Section 3 lead also to a parameterized algorithm for the MAX-SAT optimization problem which asks for the maximum number of satisfied clauses in a CNF formula. We briefly describe this extension, though we have to admit that the importance of the MAX-SAT algorithm on graphs of bounded rank-width is not as high as that of #SAT. The reason for lower applicability is that for “sparse” formula graphs (i.e. those not containing large bipartite cliques) their rank-width is bounded iff their tree-width is bounded, while for dense formula graphs the satisfiability problem is easier in general.

Theorem 4.1. *There is an algorithm that, given a CNF formula ϕ and a (t^+, t^-) -labeling parse tree of the formula graph F_ϕ , solves the MAX-SAT optimization problem of ϕ in time $\mathcal{O}(t^3 \cdot 2^{3t(t+1)/2} \cdot |\phi|)$ where $t = \max(t^+, t^-)$.*

In order to formulate this algorithm, we extend Definition 3.4 as follows. Recall $V(F_\phi) = W \cup C$ where W are the variables and C are the clauses of ϕ .

Definition 4.2. Consider a (t^+, t^-) -labeling $\tilde{F}_1 = (F_1, \text{lab}^+, \text{lab}^-)$ of a signed subgraph $F_1 \subseteq F_\phi$, and a partial assignment $\nu_1 : V(F_1) \cap W \rightarrow \{0, 1\}$. We say that ν_1 is an *assignment of defective shape* $(\Sigma^+, \Sigma^-, \Pi^+, \Pi^-)$ in \tilde{F}_1 if there exists a set $C_0 \subseteq C \cap V(F_1)$ such that ν_1 is of shape $(\Sigma^+, \Sigma^-, \Pi^+, \Pi^-)$ in $\tilde{F}_1 - C_0$. The value (the *defect*) of ν_1 with respect to this defective shape is the minimum cardinality of such C_0 .

Informally, the defect equals the number of clauses in F_1 which are unsatisfied there and not expected to be satisfied in a complete assignment in F_ϕ .

We process the parse tree T_ϕ of F_ϕ similarly to Algorithm 3.6, but this time the value of the entry $\text{Table}_z[\Sigma^+, \Sigma^-, \Pi^+, \Pi^-]$ will be equal to the minimum defect over all partial assignments in \tilde{F}_z that are of defective shape $(\Sigma^+, \Sigma^-, \Pi^+, \Pi^-)$. Formally:

Algorithm 4.3 (Theorem 4.1). Given is a CNF formula ϕ and a signed (t^+, t^-) -labeling parse tree T_ϕ of the formula graph F_ϕ .

1. We initialize all entries of $Table_z$ for $z \in V(T_\phi)$ to ∞ .
2. We process all nodes of T_ϕ in the leaves-to-root order as follows.
 - a) At a clause leaf c of T_ϕ , we set $Table_c[\emptyset, \emptyset, \Pi^+, \Pi^-] \leftarrow 1$ for all subspaces Π^+, Π^- that are both orthogonal to the label vector of $\{1\}$, and set $Table_c[\emptyset, \emptyset, \Pi^+, \Pi^-] \leftarrow 0$ otherwise.
 - b) At a variable leaf ℓ of T_ϕ , we set $Table_\ell[\langle\{1\}\rangle, \emptyset, \Pi^+, \Pi^-] \leftarrow 0$ and $Table_\ell[\emptyset, \langle\{1\}\rangle, \Pi^+, \Pi^-] \leftarrow 0$ for all pairs Π^+, Π^- .
 - c) Consider an internal node z of T_ϕ , with the left son x and the right son y such that $Table_x$ and $Table_y$ have already been computed.
 - Let the composition operators at z in the labeling parse trees T_ϕ^+, T_ϕ^- be as $\tilde{F}_z^+ = \tilde{F}_x^+ \otimes [g^+ | f_1^+, f_2^+] \tilde{F}_y^+$ and $\tilde{F}_z^- = \tilde{F}_x^- \otimes [g^- | f_1^-, f_2^-] \tilde{F}_y^-$.
 - We loop exhaustively over all indices to $Table_x, Table_y, Table_z$, i.e. over all 12-tuples of subspaces $\Sigma_x^+, \Sigma_x^-, \Pi_x^+, \Pi_x^-, \Sigma_y^+, \Sigma_y^-, \Pi_y^+, \Pi_y^-, \Sigma_z^+, \Sigma_z^-, \Pi_z^+, \Pi_z^-$ of $\text{GF}(2)^t$. If all the following are true
$$\begin{aligned} \Sigma_z^+ &= \langle f_1^+(\Sigma_x^+) \cup f_2^+(\Sigma_y^+) \rangle \text{ and } \Sigma_z^- = \langle f_1^-(\Sigma_x^-) \cup f_2^-(\Sigma_y^-) \rangle, \\ \Pi_x^+ &= \langle f_1^{+T}(\Pi_z^+) \cup g^+(\Sigma_y^+) \rangle \text{ and } \Pi_x^- = \langle f_1^{-T}(\Pi_z^-) \cup g^-(\Sigma_y^-) \rangle, \\ \Pi_y^+ &= \langle f_2^{+T}(\Pi_z^+) \cup g^{+T}(\Sigma_x^+) \rangle \text{ and } \Pi_y^- = \langle f_2^{-T}(\Pi_z^-) \cup g^{-T}(\Sigma_x^-) \rangle, \end{aligned}$$
then we let $m = Table_x[\Sigma_x^+, \Sigma_x^-, \Pi_x^+, \Pi_x^-] + Table_y[\Sigma_y^+, \Sigma_y^-, \Pi_y^+, \Pi_y^-]$. If, furthermore, $m < Table_z[\Sigma_z^+, \Sigma_z^-, \Pi_z^+, \Pi_z^-]$, then we set $Table_z[\Sigma_z^+, \Sigma_z^-, \Pi_z^+, \Pi_z^-] \leftarrow m$.
3. We find the minimum m over all the entries $Table_r[\Sigma^+, \Sigma^-, \emptyset, \emptyset]$ where r is the root of T_ϕ and Σ^+, Σ^- are arbitrary subspaces of $\text{GF}(2)^t$. An optimal solution to MAX-SAT of ϕ then has $|C| - m$ satisfied clauses.

Proof (Algorithm 4.3/Theorem 4.1, sketch). The main task is to show by means of structural induction that the algorithm correctly computes in $Table_z[\Sigma^+, \Sigma^-, \Pi^+, \Pi^-]$ the minimum defect over all partial assignments in \tilde{F}_z that are of defective shape $(\Sigma^+, \Sigma^-, \Pi^+, \Pi^-)$.

Our proof proceeds in the same way as the proof of Algorithm 3.6, with a use of the following easy claim (notation as in the referred proof):

- The defect of a partial assignment ν in \tilde{F}_z w.r.t. $(\Sigma_z^+, \Sigma_z^-, \Pi_z^+, \Pi_z^-)$ equals the sum of the defects of ν_x in \tilde{F}_x and ν_y in \tilde{F}_y w.r.t. defective shapes $(\Sigma_x^+, \Sigma_x^-, \Pi_x^+, \Pi_x^-)$ and $(\Sigma_y^+, \Sigma_y^-, \Pi_y^+, \Pi_y^-)$, respectively.

The runtime analysis follows, too. □

5 Conclusions

We have presented new FPT algorithms for the #SAT and MAX-SAT problems on formulas of bounded rank-width. Our algorithms are single-exponential in

rank-width and linear in the size of the formula, and they do not involve any “large hidden constants”. This is a significant improvement over previous results, for several reasons. In the case of tree-width this follows from the fact that rank-width is much more general than tree-width. If a graph has bounded tree-width it also has bounded rank-width, but there are classes of graphs with arbitrarily high tree-width and small rank-width (e.g. cliques, complete bipartite graphs, or distance hereditary graphs).

As for clique-width (which is bounded iff rank-width is bounded), we have obtained two significant improvements over the existing algorithms such as [7]. Firstly, rank-width can be exponentially smaller than clique-width, and therefore we obtain an exponential speed-up over the existing algorithms in the worst case. Secondly, there is an FPT algorithm for computing the rank-width of a graph (and the associated rank-decomposition) exactly, whereas in the case of clique-width we have to rely on an approximation by an exponential function of rank-width.

Finally, our paper shows that many of the recent ideas and tricks of parameterized algorithm design on graphs of bounded rank-width smoothly translate to certain SAT-related problem instances which may bring new inspiration to related research, too.

References

1. S. Arora and S. Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45:70–122, 1998.
2. H. Bodlaender. Treewidth: Algorithmic techniques and results. In *MFCS’97*, volume 1295 of *LNCS*, pages 19–36. Springer, 1997.
3. B.-M. Bui-Xuan, J. Telle, and M. Vatshelle. H-join decomposable graphs and algorithms with runtime single exponential in rankwidth. *Discrete Appl. Math.*, 158(7):809–819, 2010.
4. J. Chen and I. A. Kanj. Improved exact algorithms for MAX-SAT. *Discrete Appl. Math.*, 142(1-3):17–27, 2004.
5. D. G. Corneil and U. Rotics. On the relationship between cliquewidth and treewidth. *SIAM J. Comput.*, 34(4):825–847, 2005.
6. B. Courcelle and M. Kanté. Graph operations characterizing rank-width and balanced graph expressions. In *WG’07*, volume 4769 of *LNCS*, pages 66–75. Springer, 2007.
7. E. Fischer, J. Makowsky, and E. Ravve. Counting truth assignments of formulas of bounded tree-width or clique-width. *Discrete Appl. Math.*, 156:511–529, 2008.
8. R. Ganian and P. Hliněný. On parse trees and Myhill–Nerode-type tools for handling graphs of bounded rank-width. *Discrete Appl. Math.*, 158:851–867, 2010.
9. K. Georgiou and P. Papakonstantinou. Complexity and algorithms for well-structured k-sat instances. In *SAT’08*, pages 105–118, 2008.
10. J. Goldman and G.-C. Rota. The number of subspaces of a vector space. In W. Tutte, editor, *Recent Progress in Combinatorics*, pages 75–83. Academic Press, 1969.
11. G. Gutin, E. J. Kim, S. Szeider, and A. Yeo. Solving max-2-sat above a tight lower bound. *CoRR*, abs/0907.4573, 2009.

12. P. Hliněný and S. Oum. Finding branch-decomposition and rank-decomposition. *SIAM J. Comput.*, 38:1012–1032, 2008.
13. M. Kanté. The rank-width of directed graphs. arXiv:0709.1433v3, March 2008.
14. S. Oum. Rank-width and vertex-minors. *J. Comb. Theory Ser. B*, 95(1):79–100, 2005.
15. S. Oum. Rank-width is less than or equal to branch-width. *J. Graph Theory*, 57(3):239–244, 2008.
16. S. Oum and P. D. Seymour. Approximating clique-width and branch-width. *J. Combin. Theory Ser. B*, 96(4):514–528, 2006.
17. I. Razgon and B. O’Sullivan. Almost 2-sat is fixed-parameter tractable. *J. Comput. Syst. Sci.*, 75(8):435–450, 2009.
18. M. Samer and S. Szeider. Algorithms for propositional model counting. *J. Discrete Alg.*, 8:50–64, 2010.
19. L. G. Valiant. The complexity of computing the permanent. *Theor. Comp. Science S(2)*, pages 189–201, 1979.